arXiv:1604.01435v1 [cs.GT] 5 Apr 2016

# Approximation Algorithms for Max-Min Share Allocations of Indivisible Chores and Goods

**Haris Aziz**
NICTA and UNSW, Sydney, Australia
haris.aziz@data61.csiro.au

**Gerhard Rauchecker**
University of Regensburg
gerhard.rauchecker@wiwi.uni-regensburg.de

**Guido Schryen**
University of Regensburg
guido.schryen@wiwi.uni-regensburg.de

**Toby Walsh**
NICTA and UNSW, Sydney, Australia
toby.walsh@data61.csiro.au

## Abstract

We consider Max-min Share (MmS) allocations of items both in the case where items are goods (positive utility) and when they are chores (negative utility). We show that fair allocations of goods and chores have some fundamental connections but differences as well. We prove that like in the case for goods, an MmS allocation does not need to exist for chores and computing an MmS allocation - if it exists - is strongly NP-hard. In view of these non-existence and complexity results, we present a polynomial-time 2-approximation algorithm for MmS fairness for chores. We then introduce a new fairness concept called optimal MmS that represents the best possible allocation in terms of MmS that is guaranteed to exist. For both goods and chores, we use connections to parallel machine scheduling to give (1) an exponential-time exact algorithm and (2) a polynomial-time approximation scheme for computing an optimal MmS allocation when the number of agents is fixed.

## 1 Introduction

The fair allocation of indivisible items is a central problem in economics, computer science, and operations research [Aziz *et al.*, 2015; Brams and Taylor, 1996; Bouveret *et al.*, 2015; Lipton *et al.*, 2004]. We focus on the setting in which we have a set of $N$ agents and a set of items with each agent expressing utilities over the items. The goal is to allocate the items among the agents in a fair manner without allowing transfer of money. If all agents have positive utilities for the items, we view the items as goods. On the other hand, if all agents have negative utilities for the items, we can view the items as chores. Throughout, we assume that all agents' utilities over items are additive.

In order to identify fair allocations, one needs to formalize what fairness means. A compelling fairness concept called *Max-min Share (MmS)* was recently introduced which is weaker than traditional fairness concepts such as envy-freeness and proportionality [Bouveret and Lemaître, 2014; Budish, 2011]. An agent's MmS is the "most preferred bundle he could guarantee himself as a divider in divide-and-choose against adversarial opponents" [Budish, 2011]. The main idea is that an agent partitions the items into $N$ sets in a way that maximizes the utility of the least preferred set in the partition. The utility of the least preferred set is called the *MmS guarantee* of the agent. An allocation satisfies *MmS fairness* if each agent gets at least as much utility as her MmS guarantee. We refer to such an allocation as *MmS allocation*.[1]

Although MmS is a highly attractive fairness concept and a natural weakening of proportionality and envy-freeness [Bouveret and Lemaître, 2014; Bouveret and Lemaître, 2015], Procaccia and Wang [2014] showed that an MmS allocation of goods does not exist in general. This fact initiated research on approximate MmS allocations of goods in which each agents gets some fraction of her MmS guarantee. On the positive side, not only do MmS allocations of goods exist for most instances [Kurokawa *et al.*, 2016], but there also exists a polynomial-time algorithm that returns a 2/3-approximate MmS allocation [Procaccia and Wang, 2014; Amanatidis *et al.*, 2015]. Algorithms for computing MmS allocations of goods have been deployed and are being used for fair division [Goldman and Procaccia, 2014].

In this paper, we turn to MmS allocations for chores, a subject which has not been studied previously. Even in the more general domain of fair allocation, there is a paucity of research on chore allocation compared to goods despite there being many settings where we have chores not goods [Caragiannis *et al.*, 2012]. In general, the problem of chore allocation cannot be transformed into a problem for goods allocation [Peterson and Su, 1998].

**Contributions** We consider MmS allocation of chores for the first time and present some fundamental connections between MmS allocation of goods and chores especially when the positive utilities of the agents in the case of goods are negated to obtain a chores setting. We also show that there are differences between the two settings with no known reductions between the settings. In particular, reductions such as negating the utility values and applying an algorithm for one setting does not give an algorithm for other setting.

We show that an MmS allocation does not need to exist for chores. In view of the non-existence results, we introduce a

[1] Bouveret and Lemaître [Bouveret and Lemaître, 2014] and Budish [Budish, 2011] also formalized a fairness concept called min-Max fairness that is stronger than Max-min fairness.

new concept called *optimal MmS* for both goods and chores. An allocation is an *optimal MmS allocation* if it represents the best possible approximation of the MmS guarantee. An optimal MmS allocation has two desirable properties: (1) it always exists and (2) it satisfies MmS fairness whenever an MmS allocation exists (see Figure 1). Consequently, optimal MmS is a compelling fairness concept and a conceptual contribution of the paper.

We present bounds to quantify the gap between optimal MmS fairness and MmS fairness. For chores, we present a linear-time round-robin algorithm for this purpose that provides a 2-approximation for MmS. We show that the bound proved is *tight* for the round robin algorithm. We also show that, as in the case of goods, the computation of an MmS allocation for chores is strongly NP-hard and so is the computation of an optimal MmS allocation.

In view of the computational hardness results, we develop approximation algorithms for optimal MmS fairness. For both goods and chores, we use connections to parallel machine scheduling and some well established scheduling algorithms to derive an exponential-time exact algorithm and a PTAS (polynomial-time approximation scheme) when the number of agents is fixed. These are the first PTAS results related to MMS. As long as an MMS allocation exists (that does exist in most instances as shown analytically by Kurokawa *et al.* [2016] and experimentally by Bouveret and Lemaître [2014]), our algorithm for goods also provides a PTAS for standard MMS which in terms of approximation factor is a significant improvement over previous constant-factor approximation results [Procaccia and Wang, 2014].

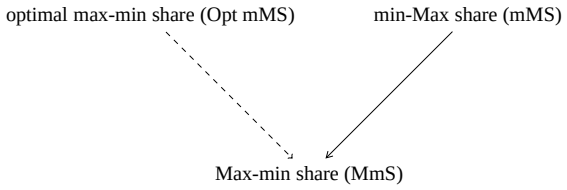


Figure 1: Relations between fairness concepts. mMS fairness implies MmS fairness. Optimal MmS fairness implies MmS fairness if an MmS allocation exists.

## 2 Related Work

In addition to the literature on MmS allocations for goods discussed in the introduction, our work is based on parallel machine scheduling theory. There is a natural connection between MmS allocations and parallel machine scheduling, which we outline later. This connection turns out to be very fruitful for both exact and approximate computations of optimal MmS allocations. We briefly introduce the concept of parallel machine scheduling in the following.

We have a set $\mathcal{M}$ of jobs and a set $[N]$ of $N$ machines. Each of the jobs has to be processed exactly once on exactly one machine without preemption. Furthermore, we have a processing time matrix $P = (p_{ij})_{i,j}$ where $p_{ij} \geq 0$ indicates how long machine $i$ requires to finish job $j$. If there are no further restrictions on the values of $P$, we deal with unrelated parallel machines. If $p_{ij} = p_{i'j}$ for all $i, i' \in [N]$ and $j \in \mathcal{M}$ then machines are considered identical.

The goal of each machine scheduling problem is to find a schedule (i.e., an ordered allocation) that optimizes a certain objective function. The problems that we focus on in this paper either minimize the time where the latest machine finishes (this is also called the makespan of a schedule) or maximize the time where the earliest machine finishes. We show that the former objective is related to MmS allocation of chores whereas the latter is related to Mms allocation of goods. An extensive overview on all important machine scheduling problems is provided by Pinedo [2012].

Graham *et al.* [1979] established a notation for machine scheduling problems where $P$ stands for identical machines, $R$ for unrelated machines, $C_{max}$ for minimizing the latest machine's finishing time, and $C_{min}$ for maximizing the earliest machine's finishing time. According to this notation, we will use the problems $P/C_{max}$, $R/C_{max}$, $P/C_{min}$, and $R/C_{min}$ in this paper. The latter problem is also equivalent to maximizing egalitarian welfare under additive utilities [Asadpour and Saberi, 2010; Bansal and Sviridenko, 2006]. All of these problems are NP-hard in the strong sense but they are well investigated and plenty of research has been conducted on approximation algorithms which we will take advantage of [Efraimidis and Spirakis, 2006; Hochbaum and Shmoys, 1987; Lenstra *et al.*, 1990; Woeginger, 1997].

## 3 Definitions

We introduce the basic notation and definitions for our approach in this section. For a set of items $\mathcal{M}$ and a number $N \in \mathbb{N}$, let $\Pi_N(\mathcal{M})$ be the set of all $N$-partitions of $\mathcal{M}$ (i.e., item allocations) and let $\mathcal{P}(\mathcal{M})$ denote the power set of $\mathcal{M}$.

**Definition 3.1.** *1. A **non-negative instance** is a tuple $(\mathcal{M}, [N], (u_i)_{i\in[N]})$ consisting of a set of items $\mathcal{M}$, a set $[N]$ of $N$ agents, and a family of additive utility functions $(u_i : \mathcal{P}(\mathcal{M}) \to \mathbb{R}_{\geq 0})_{i\in[N]}$. The set of all non-negative instances is denoted by $\mathcal{I}^+$.*

*2. A **non-positive instance** is a tuple $(\mathcal{M}, [N], (d_i)_{i\in[N]})$ consisting of a set of items $\mathcal{M}$, a set $[N]$ of $N$ agents, and a family of additive utility functions $(d_i : \mathcal{P}(\mathcal{M}) \to \mathbb{R}_{\leq 0})_{i\in[N]}$. The set of all non-positive instances is denoted by $\mathcal{I}^-$.*

*3. An **instance** is a tuple $(\mathcal{M}, [N], (v_i)_{i\in[N]})$ which is a non-negative instance or a non-positive instance. The set of all instances is denoted by $\mathcal{I}$.*

**Definition 3.2.** *For an instance $I = (\mathcal{M}, [N], (v_i)_{i\in[N]}) \in \mathcal{I}$, we define the **corresponding instance** by*

$$-I := (\mathcal{M}, [N], (-v_i)_{i\in[N]}) \in \mathcal{I}.$$

**Definition 3.3.** *Let $I = (\mathcal{M}, [N], (v_i)_{i\in[N]}) \in \mathcal{I}$ be an instance and $i \in [N]$ be an agent.*

*1. Agent $i$'s **Max-min Share (MmS) guarantee** for $I$ is defined as*

$$MmS^N_{v_i}(\mathcal{M}) := \max_{(S_1,\ldots,S_N)\in\Pi_N(\mathcal{M})} \min_{j\in[N]} v_i(S_j).$$

2. *Agent $i$'s **min-Max Share (mMS) guarantee** for $I$ is defined as*

$$mMS^N_{v_i}(\mathcal{M}) := \min_{(S_1,\dots,S_N)\in\Pi_N(\mathcal{M})} \max_{j\in[N]} v_i(S_j).$$

**Definition 3.4.** *Let $I = (\mathcal{M}, [N], (v_i)_{i\in[N]}) \in \mathcal{I}$ be an instance and $S = (S_1, \dots, S_N) \in \Pi_N(\mathcal{M})$ be an allocation.*

1. *$S$ is called an **MmS allocation** for $I$ iff $v_i(S_i) \geq MmS^N_{v_i}(\mathcal{M})$ for all agents $i \in [N]$.*
2. *$S$ is called a **perverse mMS allocation** for $I$ iff $v_i(S_i) \leq mMS^N_{v_i}(\mathcal{M})$ for all agents $i \in [N]$.*

The concept of a perverse mMS allocation seems counter-intuitive but turns out to be helpful to obtain results on MmS allocations for chores.

We can also relax the MmS fairness concept as follows.

**Definition 3.5.** *Given an instance $I = (\mathcal{M}, [N], (v_i)_{i\in[N]}) \in \mathcal{I}$ and a constant $\lambda \in \mathbb{R}$.*

1. *The **$\lambda$-max-min problem** for $I$ is about finding an allocation $(S_1, \dots, S_N) \in \Pi_N(\mathcal{M})$ with $v_i(S_i) \geq \lambda \cdot MmS^N_{v_i}(\mathcal{M})$ for all $i \in [N]$.*
2. *The **perverse $\lambda$-min-max problem** for $I$ is about finding an allocation $(S_1, \dots, S_N) \in \Pi_N(\mathcal{M})$ with $v_i(S_i) \leq \lambda \cdot mMS^N_{v_i}(\mathcal{M})$ for all $i \in [N]$.*

# 4 Properties of MmS Fairness

First, we present a fundamental connection between the allocation of chores (non-positive utilities) and goods (non-negative utilities). Later in this section, we discuss non-existence examples for MmS allocations and show that existence and non-existence examples do not transfer straightforwardly from goods to chores and vice-versa. Finally, we give a complexity result for the computation of MmS allocations for both goods and chores.

## 4.1 Fundamental Connection between Allocations of Goods and Chores

The following result shows an interesting connection between MmS and mMS when changing signs in all utility functions.

**Lemma 4.1.** *Let $(\mathcal{M}, [N], (v_i)_{i\in[N]}) \in \mathcal{I}$ be an instance. Then we have $-MmS^N_{v_i}(\mathcal{M}) = mMS^N_{-v_i}(\mathcal{M})$ for all agents $i \in [N]$.*

This leads us to the following result discussing the equivalence of MmS allocations for an instance and perverse mMS allocations for its corresponding instance.

**Proposition 4.2.** *Let $I = (\mathcal{M}, [N], (v_i)_{i\in[N]}) \in \mathcal{I}$ be an instance and let $S = (S_1, \dots, S_N) \in \Pi_N(\mathcal{M})$ be an allocation.*

1. *$S$ is an MmS allocation for $I$ if and only if $S$ is a perverse mMS allocation for the corresponding instance $-I$.*
2. *In particular, there is an MmS allocation for $I$ if and only if there is a perverse mMS allocation for the corresponding instance $-I$.*

*Proof.* For all $i \in [N]$, we have $v_i(S_i) \geq MmS^N_{v_i}(\mathcal{M}) \Leftrightarrow -v_i(S_i) \leq -MmS^N_{v_i}(\mathcal{M}) \overset{4.1}{\Leftrightarrow} -v_i(S_i) \leq mMS^N_{-v_i}(\mathcal{M})$. which proves both claims. □

This fundamental connection shows also a difference between the allocation of chores and goods since finding MmS allocations and finding perverse mMS allocations involve different objectives.

A similar statement can be made for the approximations.

**Proposition 4.3.** *Let $I = (\mathcal{M}, [N], (v_i)_{i\in[N]}) \in \mathcal{I}$ be an instance and let $S = (S_1, \dots, S_N) \in \Pi_N(\mathcal{M})$ be an allocation. Let $\lambda \in \mathbb{R}$ be arbitrary. Then $S$ is a solution of the $\lambda$-MmS problem for $I$ if and only if $S$ is a solution of the perverse $\lambda$-mMS problem for the corresponding instance $-I$.*

*Proof.* The proof is the same as the proof of Prop. 4.2 with just an additional factor $\lambda$ on the right side of each inequality. □

## 4.2 Non-existence and Complexity of MmS

Procaccia and Wang [2014] showed that an MmS allocation for goods does not necessarily exist. We construct an instance $I = (\mathcal{M}, [3], (u_i)_{i\in[3]}) \in \mathcal{I}^+$ by a subtle modification of their example to obtain an analogous result for chores. Consider a set $[3] = \{1, 2, 3\}$ of three agents and a set of twelve items (represented by pairs)

$$\mathcal{M} = \{(j, k) | j = 1, 2, 3; k = 1, 2, 3, 4\}.$$

We define matrices

$$B = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix}, \quad O = \begin{pmatrix} 17 & 25 & 12 & 1 \\ 2 & 22 & 3 & 28 \\ 11 & 0 & 21 & 23 \end{pmatrix},$$

$$E^1 = \begin{pmatrix} -3 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad E^2 = \begin{pmatrix} -3 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix},$$

$$E^3 = \begin{pmatrix} -3 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

For each agent $i \in [3]$, we define her utility function by

$$u_i : \mathcal{M} \to \mathbb{R}_{\geq 0}, \quad (j, k) \mapsto 10^6 \cdot B_{jk} + 10^3 \cdot O_{jk} + E^i_{jk}.$$

We obtain the following result by a careful adaption of the argument presented by Procaccia and Wang [2014].

**Proposition 4.4.** *There is no MmS allocation for $-I$. In particular, an MmS allocation for chores does not need to exist.*

Another interesting difference between MmS for goods and chores is the fact that existence and non-existence examples for MmS allocations cannot be simply converted into each other by just changing the signs of the utility functions.

The only difference from $I$ to the instance of the example presented by Procaccia and Wang [2014] are changed signs in the $E^i$ matrices. Let $J = (\mathcal{M}, [3], (w_i)_{i\in[3]}) \in \mathcal{I}^+$ denote their instance. We get the following interesting result.

**Proposition 4.5.** *There is an MmS allocation for $I$ but no MmS allocation for $-I$. There is no MmS allocation for $J$ but an MmS allocation for $-J$.*

Not only do MmS allocations not exist in general, computing an MmS allocation is strongly NP-hard if it exists. The reduction is straight forward from Integer Partition to an allocation instance in which each agent has the same utility function.[2]

**Proposition 4.6.** *For both goods and chores, computing an MmS allocation - if it exists - is strongly NP-hard. The problem is weakly NP-hard even for two agents.*

## 5 2-Approximation for MmS for Chores

The purpose of this section is to present a polynomial-time 2-approximation algorithm for MmS for chores. Each agent is guaranted at most twice her (non-positive) max-min share guarantee.

Through the entire section, let $I = (\mathcal{M}, [N], (u_i)_{i\in[N]}) \in \mathcal{I}^+$ be a non-negative instance. We define $u^i_{max} := \max_{j\in\mathcal{M}} u_i(j)$ for all $i \in [N]$

If we are given a chores instance $-I$, we run a round robin protocol in which in which agents come in round robin manner and are given a most preferred available item. Framed in terms of insance $I$, we consider the round-robin protocol in which agents come in round robin manner and are given an available item with the lowest utility.

**Lemma 5.1.** *Let $(S_1, \dots, S_N) \in \Pi_N(\mathcal{M})$ be the allocation obtained by the round-robin greedy protocol for $I$. Then we have $u_i(S_i) \le u_i(S_{i'}) + u^i_{max}$ for all agents $i, i' \in [N]$.*

*Proof.* The result is trivial for $i = i'$. If $i < i'$, i.e., $i$ always chooses before $i'$, then the result is also obvious because $i$ will choose her lowest valued item in every round and has to pick at most one item more than $i'$ in total (which is compensated by $u^i_{max}$).

Therefore, we can assume that $i > i'$, i.e., $i'$ picks before $i$ in every round. The protocol has exactly $K = \lceil \frac{|\mathcal{M}|}{N} \rceil$ rounds. The pick of agent $i$ ($i'$ resp.) in round $k$ is denoted by $r^i_k \in \mathcal{M}$ ($r^{i'}_k \in \mathcal{M}$ resp.). We have

$$\begin{aligned} u_i(S_{i'}) - u_i(S_i) &= \sum_{k=1}^{K} u_i(r^{i'}_k) - u_i(r^i_k) \\ &= u_i(r^{i'}_1) - u_i(r^i_1) + u_i(r^{i'}_2) - u_i(r^i_2) + \dots \\ &\dots + u_i(r^{i'}_{K-1}) - u_i(r^i_{K-1}) + u_i(r^{i'}_K) - u_i(r^i_K) \end{aligned}$$

We separate two cases. In the case that agent $i'$ has to pick an item in the last round, we have $u_i(r^i_k) \le u_i(r^{i'}_{k+1})$ for all $k = 1, \dots, K-1$ (picking rule) and therefore $u_i(S_{i'}) - u_i(S_i) \ge u_i(r^{i'}_1) - u_i(r^i_K)$. In the other case where agent $i'$ does not have to pick anymore in the last round, agent $i$ also does not have to pick since she picks after $i'$. It follows that

$$\begin{aligned} u_i(S_{i'}) - u_i(S_i) &= \sum_{k=1}^{K-1} u_i(r^{i'}_k) - u_i(r^i_k) \\ &= u_i(r^{i'}_1) - u_i(r^i_1) + u_i(r^{i'}_2) - u_i(r^i_2) + \dots \\ &\dots + u_i(r^{i'}_{K-1}) - u_i(r^i_{K-1}) \\ &\ge u_i(r^{i'}_1) - u_i(r^i_{K-1}) \end{aligned}$$

since $u_i(r^i_k) \le u_i(r^{i'}_{k+1})$ for all $k = 1, \dots, K-2$ (picking rule).

Combining both cases, this gives us $u_i(S_i) - u_i(S_{i'}) \le u_i(r^i_\alpha) - u_i(r^{i'}_1)$ with $\alpha \in \{K-1, K\}$ and finally $u_i(S_i) - u_i(S_{i'}) \le u^i_{max}$. □

**Lemma 5.2.** *Let $(S_1, \dots, S_N) \in \Pi_N(\mathcal{M})$ be the allocation obtained by the round-robin greedy protocol. Then we have*

$$u_i(S_i) \le \frac{1}{N} \cdot u_i(\mathcal{M}) + \left(1 - \frac{1}{N}\right) \cdot u^i_{max}$$

*for each agent $i \in [N]$.*

*Proof.* We have $N \cdot u_i(S_i) \overset{5.1}{\le} u_i(S_i) + \sum_{i\neq i'\in[N]} u_i(S_{i'}) + (N-1) \cdot u^i_{max} = u_i(\mathcal{M}) + (N-1) \cdot u^i_{max}$. Division by $N$ yields the result. □

**Proposition 5.3.** *The round-robin greedy allocation protocol gives an allocation $(S_1, \dots, S_N) \in \Pi_N(\mathcal{M})$ with*

$$u_i(S_i) \le \left(2 - \frac{1}{N}\right) \cdot mMS^N_{u_i}(\mathcal{M})$$

*for all $i \in [N]$. The inequality cannot be improved for general instances.*

*Proof.* By definition of the min-max share guarantee and the additivity of $u_i$, we obtain

$$\frac{1}{N} \cdot u_i(\mathcal{M}) \le mMS^N_{u_i}(\mathcal{M}) \tag{1}$$

and

$$u^i_{max} \le mMS^N_{u_i}(\mathcal{M}) \tag{2}$$

for all $i \in [N]$.

Then we have $u_i(S_i) \overset{5.2}{\le} \frac{1}{N} \cdot u_i(\mathcal{M}) + \left(1 - \frac{1}{N}\right) \cdot u^i_{max} \overset{(1)}{\le} mMS^N_{u_i}(\mathcal{M}) + \left(1 - \frac{1}{N}\right) \cdot u^i_{max} \overset{(2)}{\le} \left(2 - \frac{1}{N}\right) \cdot mMS^N_{u_i}(\mathcal{M})$ for each agent $i \in [N]$, which proves the first part of the result.

The bound cannot be improved in general. To show this, we give an example where the bound is tight. Consider a set $\mathcal{M} = (t_1, t_2, \dots, t_{(N-1)\cdot N+1})$ of $(N-1) \cdot N + 1$ items. Let the utility function $u$ be the same for all agents with $u(t_j) = \frac{1}{N}$ for all $j = 1, \dots, (N-1) \cdot N$ and $u(t_{(N-1)\cdot N+1}) = 1$.

Then we have $mMS^N_u(\mathcal{M}) = 1$ because we get a (perfectly balanced) $N$-partition by packing $N-1$ sets with exactly $N$ of the $\frac{1}{N}$-valued items and a last set consisting just of the one 1-valued item.

[2] The complexity result for goods has already been proved by Bouveret and Lemaître [2014].

Since $u$ is the same for all agents, the greedy round-robin algorithm gives agent 1 the allocation $S_1 = \{t_1, t_{N+1}, t_{2N+1} \ldots, t_{(N-1)N+1}\}$ for which

$$u(S_1) = (N-1)\cdot\frac{1}{N}+1 = 2-\frac{1}{N} = \left(2-\frac{1}{N}\right)\cdot mMS_u^N(\mathcal{M})$$

holds true. □

**Theorem 5.4.** *Let $I = (\mathcal{M}, [N], (d_i)_{i\in[N]}) \in \mathcal{I}^-$ be a non-positive instance and denote the round-robin greedy allocation for $-I$ by $(S_1, \ldots, S_N) \in \Pi_N(\mathcal{M})$. Then we have $d_i(S_i) \geq \left(2-\frac{1}{N}\right)\cdot MmS_{d_i}^N(\mathcal{M})$ for all $i \in [N]$ and the inequality cannot be improved for general instances.*

*Proof.* Follows immediately from 5.3 with 4.3. □

## 6 Optimal MmS Fairness

Before we can introduce the optimal MmS fairness concept, we have to define an instance-specific parameter.

**Definition 6.1.** *1. For a non-negative instance $I \in \mathcal{I}^+$, the **optimal MmS ratio** $\lambda^I$ is defined as the maximal $\lambda \in [0, \infty]$ for which the $\lambda$-max-min-problem for $I$ has a solution.*

*2. For a non-positive instance $I \in \mathcal{I}^-$, the **optimal MmS ratio** $\lambda^I$ is defined as the minimal $\lambda \in [0, \infty)$ for which the $\lambda$-max-min-problem for $I$ has a solution.*

Note that both the maximum and the minimum exist in this definition since for a fixed instance $I$, there is only a finite number of possible allocations. We have the following initial bounds for the optimal MmS ratio.

**Lemma 6.2.** *1. Let $I = (\mathcal{M}, [N], (u_i)_{i\in[N]}) \in \mathcal{I}^+$ be a non-negative instance. Then we have $\frac{2}{3} \leq \lambda^I \leq \infty$, with $\lambda^I = \infty$ if and only if $MmS_{u_i}^N(\mathcal{M}) = 0$ for all $i \in [N]$.*

*2. Let $I = (\mathcal{M}, [N], (d_i)_{i\in[N]}) \in \mathcal{I}^-$ be a non-positive instance. Then we have $0 \leq \lambda^I \leq 2$, and $\lambda^I = 0$ if $MmS_{d_i}^N(\mathcal{M}) = 0$ for an $i \in [N]$.*

*Proof.*

1. The inequality $\frac{2}{3} \leq \lambda^I$ is a result of the approximation algorithm presented by Procaccia and Wang [2014] while the inequality $\lambda^I \leq \infty$ is trivial. The equality $\lambda^I = \infty$ holds if and only if the $\lambda$-max-min problem for $I$ has a solution for all $\lambda \in \mathbb{R}$. This is equivalent to $MmS_{u_i}^N(\mathcal{M}) = 0$ for all $i \in [N]$.
2. The inequality $\lambda^I \geq 0$ holds per definition while $\lambda^I \leq 2$ follows from the previous section. Finally, if there is an agent $i \in [N]$ with $MmS_{d_i}^N(\mathcal{M}) = 0$ this implies $d_i \equiv 0$ and therefore allocating all items to agent $i$ gives a solution of the 0-max-min problem for $I$.

□

We do not claim that the introduced bounds of $\frac{2}{3}$ and 2 are tight. The proof of the lemma shows another difference between MmS for goods and MmS for chores. If in the case of chores, the MmS guarantee of an agent is 0, then the utility function of this agent is equal to 0. The same result does not hold true for goods.

Based on the previous notations, we define a new fairness concept called *optimal MmS*, which is a natural variant of MmS fairness.

**Definition 6.3.** *For an instance $I = (\mathcal{M}, [N], (v_i)_{i\in[N]}) \in \mathcal{I}$, an **optimal MmS allocation** for $I$ is an allocation $(S_1, \ldots, S_N) \in \Pi_N(\mathcal{M})$ with $v_i(S_i) \geq \lambda^I \cdot MmS_{v_i}^N(\mathcal{M})$ for all $i \in [N]$.*

There are two main advantages to the introduced concept. First, for each specific instance $I \in \mathcal{I}$, we can guarantee the existence of an optimal MmS allocation.[3] Second, an optimal Mms allocation is always an MmS allocation if the latter exists. Both observations follow immediately from the definitions. We will give an introductory example for optimal MmS allocations both for goods and chores in the following.

**Example 6.4.** *Define an instance $I = (\mathcal{M}, [2], (u_i)_{i\in[2]}) \in \mathcal{I}^+$ with a set $[2] = \{1, 2\}$ of two agents and a set of two items $\mathcal{M} = \{a, b\}$. We define $u_1(a) = r$, $u_1(b) = 1$, $u_2(a) = 1$, and $u_2(b) = r$ for some $r > 1$.*

1. *We have $MmS_{u_1}^2(\mathcal{M}) = MmS_{u_2}^2(\mathcal{M}) = 1$ which means that $S_1 = \{b\}$ and $S_2 = \{a\}$ is an MmS allocation for $I$ where each agent gets a total utility of 1. The optimal MmS allocation for $I$ would be $S_1 = \{a\}$ and $S_2 = \{b\}$ giving each agent a total utility of $r$. In particular, $\lambda^I = r$.*
2. *We have $MmS_{-u_1}^2(\mathcal{M}) = MmS_{-u_2}^2(\mathcal{M}) = -r$ which means that $S_1 = \{a\}$ and $S_2 = \{b\}$ is an MmS allocation for $-I$ where each agent gets a total utility of $-r$. The optimal MmS allocation for $-I$ would be $S_1 = \{b\}$ and $S_2 = \{a\}$ giving each agent a total utility of $-1$. In particular, $\lambda^{-I} = \frac{1}{r}$.*

These examples also show that each agent's ratio of the utility in an optimal MmS allocation to the utility in an MmS allocation can be arbitrarily large (for goods) or small (for chores) as $r > 1$ can be any real number. Another natural question is the worst case for the optimal MmS allocation in comparison to the MmS guarantee. This is addressed by the following definition.

**Definition 6.5.** *1. The **universal MmS ratio for goods** is defined as $\lambda^+ := \inf_{I\in\mathcal{I}^+} \lambda^I$.*

*2. The **universal MmS ratio for chores** is defined as $\lambda^- := \sup_{I\in\mathcal{I}^-} \lambda^I$.*

We can give bounds for and connections between the instance-dependent *optimal* and the instance-independent *universal* MmS ratios.

**Remark 6.6.** *Clearly, we have by definition $\lambda^+ \leq \lambda^I$ for all $I \in \mathcal{I}^+$ and $\lambda^- \geq \lambda^I$ for all $I \in \mathcal{I}^-$. Furthermore, we have:*

1. *$\frac{2}{3} \leq \lambda^+ < 1$ by Lemma 6.2 and the non-existence example presented by Procaccia and Wang [2014].*
2. *$1 < \lambda^- \leq 2$ by Prop. 4.4 and Lemma 6.2.*

We presented some properties of optimal MmS allocations. But since the complexity of computing an MmS allocation for both goods and chores - if it exists - is strongly NP-hard (Prop. 4.6), the same holds true for the computation of an optimal MmS allocation. However, we will show in the next

[3] We have to set $\infty \cdot 0 := 0$ according to Lemma 6.2.

sections that there is a PTAS for the computation of such an allocation as long as the number of agents is fixed.

## 7 Exact Algorithm and PTAS for Optimal MmS Fairness (Chores)

In this section, we develop a PTAS for finding an optimal MmS allocation for chores when the number of agents is fixed. The PTAS is based on the following exact algorithm.

**Algorithm 7.1.** *Given a non-negative instance $(\mathcal{M}, [N], (u_i)_{i\in[N]}) \in \mathcal{I}^+$, we state an algorithm consisting of the following steps.*

1. *Compute $c_i := mMS^N_{u_i}(\mathcal{M})$ for each agent $i \in [N]$.*
2. *Define new additive utility functions $u'_i : \mathcal{M} \to \mathbb{R}_{\geq 0}$ for all $i \in [N]$ by*
$$u'_i(j) := \frac{1}{c_i} \cdot u_i(j) \quad \forall j \in \mathcal{M}$$
*if $c_i > 0$ and $u'_i :\equiv 0$ if $c_i = 0$.*
3. *Solve the $R/C_{max}$ problem where each machine represents one agent and the processing times are defined as $p_{ij} := u'_i(j)$ for all $i \in [N]$ and $j \in \mathcal{M}$. Denote the optimal objective function value by $\lambda^*$ and the corresponding allocation by $S^* = (S^*_1, \ldots, S^*_N) \in \Pi_N(\mathcal{M})$.*

**Proposition 7.2.** *Given a non-positive instance $I = (\mathcal{M}, [N], (d_i)_{i\in[N]}) \in \mathcal{I}^-$. Execute algorithm 7.1 for the corresponding instance $(\mathcal{M}, [N], (u_i)_{i\in[N]}) := -I \in \mathcal{I}^+$. Then we have $\lambda^* = \lambda^I$ and $S^*$ is an optimal MmS allocation for $I$.*

*Proof.* We will show that $\lambda^*$ is the minimal $\lambda \in [0, \infty)$ for which a solution to the perverse $\lambda$-min-max problem for $-I$ exists and $S^*$ is a corresponding solution. The claim follows then with Prop. 4.3.

If $c_i = 0$ (and therefore $u_i \equiv 0$) for an agent $i \in [N]$, we have $\lambda^* = 0$ (because we can give all items to $i$) and there is nothing to show. Let us now assume $c_i > 0$ for all $i \in [N]$. $\lambda^*$ is by definition the minimal $\lambda \geq 0$ for which an allocation $(S_1, \ldots, S_N) \in \Pi_N(\mathcal{M})$ with $u'_i(S_i) \leq \lambda$ for all $i \in [N]$ exists. But this condition is equivalent to $u_i(S_i) \leq \lambda \cdot mMS^N_{u_i}(\mathcal{M})$ for all $i \in [N]$. To sum up, $\lambda^*$ is the minimal $\lambda \geq 0$ for which a solution of the perverse $\lambda$-min-max problem for $-I$ exists and $S^*$ is a corresponding solution. □

There are two steps in algorithm 7.1 that are exponential in time. First, each computation of $mMS^N_{u_i}(\mathcal{M})$ may require exponential time and second, finding an optimal solution to $R/C_{max}$ may require exponential time. The computation of $mMS^N_{u_i}(\mathcal{M})$ for an agent $i \in [N]$ is equivalent to the computation of a job partition that minimizes the makespan on $N$ identical parallel machines ($P/C_{max}$) where the processing time of a job $j \in \mathcal{M}$ is defined as $p_j := u_i(j)$.

Hochbaum and Shmoys [1987] present a PTAS for $P/C_{max}$ and Lenstra *et al.* [1990] present a PTAS for $R_N/C_{max}$ (which means that the number of agents is fixed to $N$). This implies that we can run the following algorithm in polynomial time when the number of agents is fixed and therefore, it will be suitable to formulate a PTAS for the computation of an optimal MmS allocation when the number of agents is fixed.

**Algorithm 7.3.** *Given a non-negative instance $I = (\mathcal{M}, [N], (u_i)_{i\in[N]}) \in \mathcal{I}^+$ and an $\varepsilon > 0$, we state an algorithm consisting of the following steps.*

1. *Select $\alpha > 1$ and $\beta > 1$ with $\alpha\beta < 1 + \varepsilon$.*
2. *Compute $c_i$ with $mMS^N_{u_i}(\mathcal{M}) \leq c_i \leq \alpha \cdot mMS^N_{u_i}(\mathcal{M})$ for each agent $i \in [N]$.*
3. *Define new additive utility functions $u'_i : \mathcal{M} \to \mathbb{R}_{\geq 0}$ for all $i \in [N]$ by*
$$u'_i(j) := \frac{1}{c_i} \cdot u_i(j) \quad \forall j \in \mathcal{M}$$
*if $c_i > 0$ and $u'_i :\equiv 0$ if $c_i = 0$.*
4. *Consider the corresponding $R/C_{max}$ problem where the processing times are defined as $p_{ij} := u'_i(j)$ for all $i \in [N]$ and $j \in \mathcal{M}$. Denote the optimal objective function value by $\lambda^*$. Compute an approximate solution $S^\varepsilon = (S^\varepsilon_1, \ldots, S^\varepsilon_N) \in \Pi_N(\mathcal{M})$ with $u'_i(S^\varepsilon_i) \leq \beta\lambda^*$ for all $i \in [N]$.*

**Theorem 7.4.** *Let a non-positive instance $I = (\mathcal{M}, [N], (d_i)_{i\in[N]}) \in \mathcal{I}^-$, $\lambda \geq 0$, and $\varepsilon > 0$ be given. Assume that a solution of the $\lambda$-max-min problem for $I$ exists and execute Algorithm 7.3 for the pair $(-I, \epsilon)$. Then $S^\varepsilon$ is a solution of the $((1 + \varepsilon) \cdot \lambda)$-max-min problem for $I$.*

*Proof.* Let $u_i$ refer to $-d_i$ for all $i \in [N]$. Since a solution of the $\lambda$-max-min problem for $I$ exists, we can conclude by 4.3 that a solution of the perverse $\lambda$-min-max problem for $-I$ exists.

This implies the existence of $(S_1, \ldots, S_N) \in \Pi_N(\mathcal{M})$ with $u_i(S_i) \leq \lambda \cdot mMS^N_{u_i}(\mathcal{M}) \leq \lambda \cdot c_i$ for all $i \in [N]$. From this we have $u'_i(S_i) \leq \lambda$ for all $i \in [N]$ (note that $c_i = 0$ implies $u'_i \equiv 0$) and we can conclude $\lambda^* \leq \lambda$.

Define $[N]_{>0} := \{i \in [N] | c_i > 0\}$. This gives us

$$\max_{i\in[N]_{>0}} \frac{u_i(S^\varepsilon_i)}{c_i} = \max_{i\in[N]_{>0}} u'_i(S^\varepsilon_i) \leq \beta\lambda^* \leq \beta\lambda.$$

Since $c_i \leq \alpha \cdot mMS^N_{u_i}(\mathcal{M})$, this leads us to

$$\max_{i\in[N]_{>0}} \frac{u_i(S^\varepsilon_i)}{mMS^N_{u_i}(\mathcal{M})} \leq \alpha\beta\lambda \leq (1 + \varepsilon) \cdot \lambda$$

which is equivalent to

$$u_i(S^\varepsilon_i) \leq (1 + \varepsilon) \cdot \lambda \cdot mMS^N_{u_i}(\mathcal{M})$$

for all $i \in [N]_{>0}$ and the same is true for all $i \in [N] \backslash [N]_{>0}$ (since this means $u_i \equiv 0$).

This proves that $S^\varepsilon$ is a solution of the perverse $((1+\varepsilon)\cdot\lambda)$-min-max problem for $-I$. The result follows now by 4.3. □

Since Algorithm 7.3 runs in polynomial time when the number of agents is fixed, this general result gives us immediately the following important corollary.

**Corollary 7.5.** *Let the number of agents be fixed to $N$ and let $I \in \mathcal{I}^-$ be a non-positive instance with $N$ agents.*

1. *Executing algorithm 7.3 for $-I$ and $\varepsilon > 0$ gives a PTAS for the computation of an optimal MmS allocation for $I$.*
2. *If an MmS allocation for $I$ exists, then executing algorithm 7.3 for $-I$ and $\varepsilon > 0$ gives a PTAS for the computation of an MmS allocation for $I$.*

*Proof.* Set $\lambda = \lambda^I$ in 7.4 for the first result. The second statement follows since each optimal MmS allocation for $I$ is also an MmS allocation for $I$ if the latter exists. □

This is a strong result since it gives a PTAS for the computation of an optimal MmS allocation of a given chore instance, no matter if an MmS allocation exists or not (assuming the number of agents is fixed). In addition, if an MmS allocation for a chore instance exists, we have a PTAS (assuming the number of agent is fixed) to compute an MmS allocation. An analogous result for goods will be obtained in the next section.

## 8 Exact Algorithm and PTAS for Optimal MmS Fairness (Goods)

In this section, we present a PTAS for finding an optimal MmS allocation for goods when the number of agents is fixed. The PTAS is based on the following exact algorithm. The techniques for the proofs are basically the same as in the previous section.

**Algorithm 8.1.** *Given a non-negative instance $(\mathcal{M}, [N], (u_i)_{i\in[N]}) \in \mathcal{I}^+$, we state an algorithm consisting of the following steps.*

1. *Compute $c_i := MmS^N_{u_i}(\mathcal{M})$ for each agent $i \in [N]$ and set*
$$[N]_{>0} := \{i \in [N] | c_i > 0\}.$$
2. *If $[N]_{>0} = \emptyset$, set $\lambda^* := \infty$ and choose an arbitrary allocation $S^* \in \Pi_N(\mathcal{M})$. Terminate the algorithm.*
3. *Define new additive utility functions*
$$u_i' : \mathcal{M} \to \mathbb{R}_{\geq 0}, \quad j \mapsto \frac{1}{c_i} \cdot u_i(j)$$
*for all $i \in [N]_{>0}$.*
4. *Solve the $R/C_{min}$ problem on $|[N]_{>0}|$ machines where each machine represents one agent $i \in [N]_{>0}$ and the processing times are defined as $p_{ij} := u_i'(j)$ for all $i \in [N]_{>0}$ and $j \in \mathcal{M}$. Denote the optimal objective function value by $\lambda^*$ and the corresponding allocation by $(S_i^*)_{i\in[N]_{>0}} \in \Pi_{|[N]_{>0}|}(\mathcal{M})$. Set $S_i^* := \emptyset$ for all $i \in [N]\backslash[N]_{>0}$ and $S^* := (S_1^*, \ldots, S_N^*) \in \Pi_N(\mathcal{M})$.*

Note that this algorithm aims at finding an optimal MmS allocation by maximizing egalitarian welfare according to the new utility functions $u_i'$.

**Proposition 8.2.** *Execute Algorithm 8.1 for a non-negative instance $I = (\mathcal{M}, [N], (u_i)_{i\in[N]}) \in \mathcal{I}^+$. Then we have $\lambda^* = \lambda^I$ and $S^*$ is an optimal MmS allocation for $I$.*

There are again two steps in Algorithm 8.1 that are exponential in time. First, each computation of $MmS^N_{u_i}(\mathcal{M})$ may require exponential time and second, finding an optimal solution to $R/C_{min}$ may require exponential time. The computation of $MmS^N_{u_i}(\mathcal{M})$ for an agent $i \in [N]$ is equivalent to the computation of a job partition that maximizes the minimum finishing time on $N$ identical parallel machines ($P/C_{min}$) where the processing time of a job $j \in \mathcal{M}$ is defined as $p_j := u_i(j)$.

Woeginger [1997] present a PTAS for $P/C_{min}$ and Efraimidis and Spirakis [2006] present a PTAS for $R_N/C_{min}$ (which means that the number of agents is fixed to $N$).[4] This implies that we can run the following algorithm in polynomial time when the number of agents is fixed.

**Algorithm 8.3.** *Given a non-negative instance $I = (\mathcal{M}, [N], (u_i)_{i\in[N]}) \in \mathcal{I}^+$ and an $\varepsilon > 0$, we state an algorithm consisting of the following steps.*

1. *Select $0 < \alpha < 1$ and $0 < \beta < 1$ with $\alpha\beta > 1 - \varepsilon$.*
2. *Compute $c_i$ with $\alpha \cdot MmS^N_{u_i}(\mathcal{M}) \leq c_i \leq MmS^N_{u_i}(\mathcal{M})$ for each agent $i \in [N]$. Define a set $[N]_{>0} := \{i \in [N] | c_i > 0\}$.*
3. *If $[N]_{>0} = \emptyset$, set $\lambda^* := \infty$ and choose an arbitrary allocation $S^\varepsilon \in \Pi_N(\mathcal{M})$. Terminate the algorithm.*
4. *Define new additive utility functions*
$$u_i' : \mathcal{M} \to \mathbb{R}_{\geq 0}, \quad j \mapsto \frac{1}{c_i} \cdot u_i(j)$$
*for all $i \in [N]_{>0}$.*
5. *Consider the corresponding $R/C_{min}$ problem on $|N_{>0}|$ machines where each machine represents one agent $i \in [N]_{>0}$ and the processing times are defined as $p_{ij} := u_i'(j)$ for all $i \in [N]_{>0}$ and $j \in \mathcal{M}$. Denote the optimal objective function value by $\lambda^*$. Compute an approximate solution $(S_i^\varepsilon)_{i\in[N]_{>0}} \in \Pi_{|[N]_{>0}|}(\mathcal{M})$ with $u_i'(S_i^\varepsilon) \geq \beta\lambda^*$ for all $i \in [N]_{>0}$. Set $S_i^\varepsilon := \emptyset$ for all $i \in [N]\backslash[N]_{>0}$ and $S^\epsilon := (S_1, \ldots, S_N) \in \Pi_N(\mathcal{M})$.*

**Theorem 8.4.** *Let a non-negative instance $I = (\mathcal{M}, [N], (u_i)_{i\in[N]}) \in \mathcal{I}^+$, $\lambda \geq 0$, and $\varepsilon > 0$ be given. Assume that a solution of the $\lambda$-max-min problem for $I$ exists and execute Algorithm 8.3 for the pair $(I, \varepsilon)$. Then $S^\epsilon$ is a solution of the $((1-\varepsilon)\cdot\lambda)$-max-min problem for $I$.*

Since Algorithm 8.3 runs in polynomial time when the number of agents is fixed, this general result gives us the following important corollary.

**Corollary 8.5.** *Let the number of agents be fixed to $N$ and let $I \in \mathcal{I}^+$ be a non-negative instance with $N$ agents.*

1. *Algorithm 8.3 gives a PTAS for the computation of an optimal MmS allocation for $I$.*
2. *If an MmS allocation for $I$ exists, then Algorithm 8.3 gives a PTAS for the computation of an MmS allocation for $I$.*

## 9 Conclusions

We initiated work on MmS allocation of chores, proposed a new fairness concept called optimal MmS, and presented interesting connections and differences between fair allocation of goods and chores. For a fixed number of agents, we proposed compelling approximation algorithms for fair allocation with respect to MmS for both chores and goods. For goods, we provided a connection between computation of MmS allocations and egalitarian welfare maximizing allocations. There is much potential for further work on the fair allocation of chores as this has largely been overlooked in contrast to the case of goods. One interesting direction are problems containing a mixture of goods and chores. For such

[4] A PTAS for each fixed number of agents implies also a PTAS for each bounded number of at most $N$ agents.

settings, an approximate MmS allocation will require approximation guarantees on either side depending on whether the MmS guarantee is positive or negative.

# Appendix

## A Proof of Proposition 4.4

The matrix $O$ gets the following labeling:

$$\begin{pmatrix} {}^{\alpha}17^{1}_{+} & {}^{\alpha}25^{1}_{-} & {}^{\beta}12^{1}_{+} & {}^{\gamma}1^{1}_{*} \\ {}^{\alpha}2^{2}_{-} & {}^{\beta}22^{2}_{*} & {}^{\gamma}3^{2}_{+} & {}^{\gamma}28^{2}_{-} \\ {}^{\alpha}11^{3}_{*} & {}^{\beta}0^{3}_{-} & {}^{\beta}21^{3}_{*} & {}^{\gamma}23^{3}_{+} \end{pmatrix}$$

As Procaccia and Wang [2014] point out, each 4 elements in $O$ sum up to 55 if and only if they have a common label. Suppose agent 1 divides the items along the labels $1, 2, 3$, agent 2 divides the items along the labels $\alpha, \beta, \gamma$ and agent 3 divides the items along the labels $+, -, *$, then every agent has perfectly balanced bundles where the utility is always $4,055,000$. Therefore, we can conclude

$$\begin{aligned} mMS^3_{u_1}(\mathcal{M}) = mMS^3_{u_2}(\mathcal{M}) = mMS^3_{u_3}(\mathcal{M}) \\ = 4,055,000. \end{aligned}$$

Consider an arbitrary allocation $(S_1, S_2, S_3) \in \Pi_3(\mathcal{M})$. We will show that there is always an agent $i \in [3]$ with $u_i(S_i) > 4,055,000 = mMS^3_{u_i}(\mathcal{M})$.

If there is an agent $i \in [3]$ which receives at least five items, we have $u_i(S_i) > 5,000,000 > 4,055,000$. Consequently, we can focus on the case where each agent receives exactly four items. If there is an agent $i \in [3]$ which receives items with a sum of more than 55 in the O-matrix, we have $u_i(S_i) > 4,055,096 > 4,055,000$. Consequently, we can focus on the case where all agents get a bundle with common labels (i.e., with a sum of exactly 55 in the $O$-matrix).

If the items are divided along the labels $1, 2, 3$ then agent 2 or agent 3 receive items labeled by 2 or 3 giving them a utility of $4,055,001$. If the items are divided along the labels $\alpha, \beta, \gamma$ then agent 1 or agent 3 receive items labeled by $\beta$ or $\gamma$ giving them a utility of at least $4,055,001$. If the items are divided along the labels $+, -, *$ then agent 1 or agent 2 receive items labeled by $-$ or $*$ giving them a utility of at least $4,055,001$.

We showed that for an arbitrary allocation $(S_1, S_2, S_3) \in \Pi_3(\mathcal{M})$ there is always an agent $i \in [3]$ with $u_i(S_i) > 4,055,000 = mMS^3_{u_i}(\mathcal{M})$ which means that there is no perverse mMS allocation for the instance $I$. The result follows from 4.2. □

## B Proof of Proposition 4.5

We define

$$S_i := \{(i,k) | k = 1, 2, 3, 4\}$$

for all $i \in [3]$.

As pointed out in Appendix A, we can achieve a perfectly balanced partition with

$$MmS^3_{u_i}(\mathcal{M}) = mMS^3_{u_i}(\mathcal{M}) = 4,055,000$$

for all $i \in [3]$. Furthermore, we have $u_1(S_1) = 4,055,000$ and $u_2(S_2) = u_3(S_3) = 4,055,001$, which means

$$u_i(S_i) \geq MmS^3_{u_i}(\mathcal{M})$$

for all $i \in [3]$ and therefore $(S_1, S_2, S_3)$ is an MmS allocation for $\mathcal{I}$. The non-existence of an MmS allocation for $-I$ was shown in 4.4.

The utility functions $w_i$ for each $i \in [3]$ are defined by

$$w_i((j,k)) := 10^6 \cdot B_{jk} + 10^3 \cdot O_{jk} - E^i_{jk}.$$

This implies $w_1(S_1) = 4,055,000$ and $w_2(S_2) = w_3(S_3) = 4,054,999$. As Procaccia and Wang [2014] point out, we have $mMS^3_{w_i}(\mathcal{M}) = 4,055,000$ for all $i \in [3]$, which means that $(S_1, S_2, S_3)$ is a perverse mMS allocation for $J$ and therefore an MmS allocation for $-J$ by 4.2. The non-existence of an MmS allocation for $J$ was shown by Procaccia and Wang [2014]. □

## C Proof of Proposition 4.6

We show the strong NP-hardness by a reduction from 3-partition. We consider a setting with numbers $N, W \in \mathbb{N}$, a set $\mathcal{M} = \{1, 2, \ldots, 3N\}$ of $3N$ elements, and an additive valuation function $u : \mathcal{M} \to \mathbb{R}_{>0}$ such that $\frac{W}{4} < u(j) < \frac{W}{2}$ for all $j = 1, \ldots, 3N$ and $\sum_{j=1}^{3N} u(j) = NW$. Question: Can $\mathcal{M}$ can be partitioned into $N$ disjoint subsets where the total valuation of the elements in each subset is $W$? This decision problem is a strongly NP-complete restricted version of the 3-partition problem [Garey and Johnson, 1979].

The computation of an MmS (resp. perverse mMS) allocation for the corresponding instance $(\mathcal{M}, [N], (u)_{i\in[N]})$ where each agent has the same utility function $u$ is equivalent to the computation of an MmS (resp. perverse mMS) partition for any single agent.

But this would answer the above mentioned strongly NP-complete decision problem and therefore the computation of an MmS allocation for both goods and chores (by 4.2) - if it even exists - is strongly NP-hard.

Let us now fix the number of agents to 2. We consider a number $M \in \mathbb{N}$, a set $\{1, 2, \ldots, M\}$ and an additive valuation function $u : \mathcal{M} \to \mathbb{R}_{>0}$. Question: Can $\mathcal{M}$ can be partitioned into 2 disjoint subsets where the total valuation of the elements in both subsets is the same? This is a general instance of the integer partition problem which is NP-complete [Garey and Johnson, 1979].

The computation of an MmS (resp. perverse mMS) allocation for the corresponding instance $(\mathcal{M}, [2], (u)_{i\in[2]})$ where both agents have the same utility function $u$ is equivalent to the computation of an MmS (resp. perverse mMS) partition for any single agent.

But this would answer the NP-complete integer partition decision problem and therefore the computation of an MmS allocation for both goods and chores (by 4.2) - if it even exists - is NP-hard for $N = 2$. □

## D Proof of Proposition 8.2

If $c_i = 0$ for all $i \in [N]$ then we have $\lambda^* = \lambda^I = \infty$ by 6.2. So let us assume that $c_i > 0$ for at least one $i \in [N]$. By construction, $\lambda^*$ is the maximal $\lambda \geq 0$ for which an allocation $(S_i)_{i\in[N]_{>0}} \in \Pi_{|N_{>0}|}(\mathcal{M})$ with $u'_i(S_i) \geq \lambda$ for all $i \in [N]_{>0}$ exists. If $c_i = 0$ for an agent $i \in [N]$ then even

$$u_i(\emptyset) \geq \lambda \cdot MmS^N_{u_i}(\mathcal{M})$$

holds true for all $\lambda \in [0, \infty)$.

In summary, $\lambda^*$ is the maximal $\lambda \geq 0$ for which an allocation $(S_i)_{i\in[N]} \in \Pi_N(\mathcal{M})$ with $u_i(S_i) \geq \lambda \cdot MmS^N_{u_i}(\mathcal{M})$ for all $i \in [N]$ exists and therefore $\lambda^* = \lambda^I$. □

## E Proof of theorem 8.4

If $c_i = 0$ for all $i \in [N]$, there is nothing to show. So let us assume $[N]_{>0} \neq \emptyset$. Since we know that a solution of the $\lambda$-max-min problem for $I$ exists, there must be an allocation $(S_1, \ldots, S_N) \in \Pi_N(\mathcal{M})$ with $u_i(S_i) \geq \lambda \cdot MmS^N_{u_i}(\mathcal{M}) \geq \lambda \cdot c_i$ for all $i \in [N]$. From this, we have $u'_i(S_i) \geq \lambda$ for all $i \in [N]_{>0}$, which implies $\lambda \leq \lambda^*$.

Therefore, we have

$$\min_{i\in[N]_{>0}} \frac{u_i(S^\varepsilon_i)}{c_i} = \min_{i\in[N]_{>0}} u'_i(S^\varepsilon_i) \geq \beta\lambda^* \geq \beta\lambda.$$

Since $c_i \geq \alpha \cdot mMS^N_{u_i}(\mathcal{M})$ for all $i \in [N]$, this leads us to

$$\min_{i\in[N]_{>0}} \frac{u_i(S^\varepsilon_i)}{MmS^N_{u_i}(\mathcal{M})} \geq \alpha\beta\lambda \geq (1-\varepsilon)\cdot\lambda.$$

But this means nothing else than

$$u_i(S^\varepsilon_i) \geq (1-\varepsilon)\cdot\lambda\cdot MmS^N_{u_i}(\mathcal{M})$$

for all $i \in [N]_{>0}$ which holds clearly also for $i \in [N]\backslash[N]_{>0}$. □

## References


[Amanatidis *et al.*, 2015] G. Amanatidis, E. Markakis, A. Nikzad, and A. Saberi. Approximation algorithms for computing maximin share allocations. In *Proceedings of the 35th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 39–51, 2015.

[Asadpour and Saberi, 2010] A. Asadpour and A. Saberi. An approximation algorithm for max-min fair allocation of indivisible goods. *SIAM Journal on Computing*, 39(7):2970–2989, 2010.

[Aziz *et al.*, 2015] H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92, 2015.

[Bansal and Sviridenko, 2006] N. Bansal and M. Sviridenko. The santa claus problem. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 31–40. ACM Press, 2006.

[Bouveret and Lemaître, 2014] S. Bouveret and M. Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1321–1328. IFAAMAS, 2014.

[Bouveret and Lemaître, 2015] S. Bouveret and M. Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 2015.

[Bouveret *et al.*, 2015] S. Bouveret, Y. Chevaleyre, and N. Maudet. Fair allocation of indivisible goods. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 12. Cambridge University Press, 2015.

[Brams and Taylor, 1996] S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.

[Budish, 2011] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

[Caragiannis *et al.*, 2012] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. The efficiency of fair division. *Theory of Computing Systems*, 50(4):589–610, 2012.

[Efraimidis and Spirakis, 2006] P. S. Efraimidis and P. G. Spirakis. Approximation schemes for scheduling and covering on unrelated machines. *Theoretical Computer Science*, 359(1):400–417, 2006.

[Garey and Johnson, 1979] Michael R Garey and David S Johnson. *Computers and intractability: a guide to the theory of NP-completeness. 1979*. Freeman & Co, 1979.

[Goldman and Procaccia, 2014] J. Goldman and A. D. Procaccia. Spliddit: Unleashing fair division algorithms. *ACM SIGecom Exchanges*, 13(2):41–46, 2014.

[Graham *et al.*, 1979] R. L. Graham, E. L. Lawler, J. K. Lenstra, and A. H. G. Rinnooy Kan. Optimization and approximation in deterministic sequencing and scheduling: A survey. *Annals of Discrete Mathematics*, 5:287–326, 1979.

[Hochbaum and Shmoys, 1987] D. S. Hochbaum and D. B. Shmoys. Using dual approximation algorithms for scheduling problems theoretical and practical results. *Journal of the ACM (JACM)*, 34(1):144–162, 1987.

[Kurokawa *et al.*, 2016] D. Kurokawa, A. D. Procaccia, and J. Wang. When can the maximin share guarantee be guaranteed? In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, 2016.

[Lenstra *et al.*, 1990] J. K. Lenstra, D. B. Shmoys, and E. Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Mathematical programming*, 46(1-3):259–271, 1990.

[Lipton *et al.*, 2004] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (ACM-EC)*, pages 125–131. ACM Press, 2004.

[Peterson and Su, 1998] E. Peterson and F. E. Su. Exact procedures for envy-free chore division. 1998.

[Pinedo, 2012] M. L. Pinedo. *Scheduling: Theory, Algorithms, and Systems*, volume 4. Springer, 2012.

[Procaccia and Wang, 2014] A. D. Procaccia and J. Wang. Fair enough: Guaranteeing approximate maximin shares.

In *Proceedings of the 15th ACM Conference on Economics and Computation (ACM-EC)*, pages 675–692. ACM Press, 2014.

[Woeginger, 1997] G. J. Woeginger. A polynomial-time approximation scheme for maximizing the minimum machine completion time. *Operations Research Letters*, 20(4):149–154, 1997.